Seminar Hora Informaticae **Institute of Computer Science,Prague**

Tuesday, January 30, 2024, 14.00 – 15.30 (2 – 3:30 PM) CET

Meeting Room 318, Address: Pod Vodárenskou věží 2, Prague 8

ZOOM Meeting ID: 954 7823 4977 , Passcode: 712564

ZOOM: https://cesnet.zoom.us/j/95478234977?pwd=dXoyekFHbDJ0MkNrTjVVS3F2STZqUT09

Martin Berger (Montanarius Ltd & University of Sussex):

GPU-acceleration of program synthesis

GPUs are the work-horses of high-performance computing. The acceleration they provide to applications compatible with their programming paradigm can surpass CPU performance by several orders of magnitude, as notably evidenced by the advancements in deep learning. A significant spectrum of applications, especially within automated reasoning has yet to reap the benefits of GPU acceleration. In order for an application to be “GPU-friendly”, it needs to have high parallelism, minimal data-dependent branching, and predictable data movement with substantial data locality. Current automated reasoning algorithms are predominantly branching-intensive and appear sequential in nature, but it is unclear whether they are inherently sequential, or can be adapted to GPUs.

We identify program synthesis as a candidate for GPU-acceleration. Program synthesis is an umbrella term for the algorithmic generation of programs (and similar formal objects, like logical formulae) from specifications. We conjecture that program synthesis is especially suitable for GPU-acceleration because it is embarrassingly parallel: most program synthesis involves “generate-and-check” phases where vast numbers of candidate solutions are first generated, and then checked if they meet the target specification. Synthesis stops as soon as a suitable candidate is found. This can be done in parallel with only moderate amounts of synchronisation, making program synthesis

embarrassingly parallel and potentially GPU friendly. The (typically) exponential growth in potential synthesis candidates will effortlessly saturate the compute cores of any future processor.

In order to test our hypothesis, we implement two classic machine learning problems using program synthesis on a GPU (regular expression inference from examples and linear-temporal logic learning, also from examples). We observe several orders of magnitude speedup, and an ability to handle much larger instances.

References:

https://www.sussex.ac.uk/foss/publications/berger

Martin Bergen (https://martinfriedrichberger.net/) is a Senior Lecturer in Foundations of Computation (Informatics) School of Engineering and Informatics, university of Sussex.

HORA INFORMATICAE (meaning: TIME FOR INFORMATICS) is a broad-spectrum scientific seminar devoted to all core areas of computer science and its interdisciplinary interfaces with other sciences and applied domains. Original contributions addressing classical and emerging topics are welcome. Founded by Jiří Wiedermann, the seminar is running since 1994 at the Institute of Computer Science of the Czech Academy of Sciences in Prague.

https://www.cs.cas.cz/horainf